

Article

Seasonal Precipitation Provides Modest Incremental Information for Retrospective Estimation of Active-Layer Thickness at Monitored Sites Along the Qinghai–Tibet Engineering Corridor, China

Qingsong Du ^{1,2,3,4,*} , Fei Wang ^{1,2,5} , Guoyu Li ^{1,2,5} , Dun Chen ^{1,2,5}  and Shunshun Qi ^{1,2} 

- ¹ State Key Laboratory of Cryospheric Science and Frozen Soil Engineering, Northwest Institute of Eco-Environment and Resources, Chinese Academy of Sciences, Lanzhou 730000, China; wangfei@nieer.ac.cn (F.W.); guoyuli@lzb.ac.cn (G.L.); chendun@lzb.ac.cn (D.C.); qishunshun@nieer.ac.cn (S.Q.)
- ² Da Xing'anling Observation and Research Station of Frozen-Ground Engineering and Environment, Northwest Institute of Eco-Environment and Resources, Chinese Academy of Sciences, Da Xing'anling 165000, China
- ³ Department of Earth and Atmospheric Sciences, University of Alberta, Edmonton, AB T6G 2E3, Canada
- ⁴ Qinghai-Beiluhe Plateau Frozen Soil Engineering Safety National Observation and Research Station, Northwest Institute of Eco-Environment and Resources, Chinese Academy of Sciences, Lanzhou 730000, China
- ⁵ School of Engineering Science, University of Chinese Academy of Sciences, Beijing 100049, China
- * Correspondence: xbdqs@lzb.ac.cn

Abstract

Active-layer thickness (ALT) integrates atmospheric forcing with local surface and sub-surface hydrothermal conditions, but the incremental predictive value of precipitation after temperature and site memory have been considered to remain uncertain. We combined annual ALT observations from 54 boreholes along the Xidatan–Anduo section of the Qinghai–Tibet Engineering Corridor (2001–2020) with monthly Third Pole Meteorological Forcing Dataset data (1986–2020). A rolling-origin design evaluated nine held-out years (2012–2020) after accounting for site fixed effects, a quadratic temporal trend, previous-year ALT, and an air-temperature lag selected from each training set. Because the principal predictors include target-year precipitation, the analysis represents retrospective or end-of-season annual estimation rather than a lead-time forecast. Among 16 seasonal-lag candidates, the post-screening best fixed model used June–August precipitation averaged over the target year and the three preceding years; it reduced RMSE from 0.203491 to 0.199312 m ($\Delta\text{RMSE} = 0.004179$ m; 2.05%). A selection-aware maximum statistic remained supported under site-specific circular shifts ($p = 0.001$), whereas a common-shift sensitivity that preserved synchronous cross-site climate structure was inconclusive. A pipeline that repeated precipitation-window selection using training data only produced a smaller gain ($\Delta\text{RMSE} = 0.001645$ m; 0.81%; $p = 0.004$). Ridge regression and random forest showed similarly small paired gains, but histogram gradient boosting did not. A continuous mean annual ground temperature (MAGT) interaction did not improve external performance, strict spatial-block $\times$ time holdout showed no transferable precipitation gain, and a predictor restricted to precipitation from years $t - 1$ to $t - 4$ did not improve RMSE. The fixed-window gain was substantially smaller than the source-reported ± 0.05 m uncertainty of an individual ALT estimate. Seasonal precipitation should therefore be treated as an auxiliary covariate for retrospective assessment at monitored sites, not as a stand-alone forecast, an ungauged-site hazard model, or evidence of maintenance-cost savings.



Academic Editor: Bingxiang Yuan

Received: 22 June 2026

Revised: 14 July 2026

Accepted: 28 July 2026

Published: 30 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)

[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: permafrost; active-layer thickness; seasonal precipitation; retrospective estimation; rolling-origin evaluation; temporal null; Qinghai–Tibet Engineering Corridor

1. Introduction

Permafrost on the Tibetan Plateau is warming and degrading, with consequences that include active-layer deepening, rising ground temperatures, and declining frozen-ground thermal stability [1–4]. Because these changes couple thermal, hydrological, ecological, and infrastructure questions, their evaluation increasingly requires cross-disciplinary evidence [5]. They are especially important along the Qinghai–Tibet Engineering Corridor (QTEC), which contains the Qinghai–Tibet Highway (QTH), Qinghai–Tibet Railway (QTR), pipelines, transmission lines, and other infrastructure [6]. Permafrost degradation, thaw of ice-rich soil, and development of thaw bulbs can contribute to settlement, embankment deformation, cracking, drainage disturbance, and loss of serviceability [7–12].

Published evidence also shows that the economic context is substantial. A regional engineering review reported RMB 0.35–0.47 billion in direct losses associated with frozen-ground environmental damage before 2005 [12]. More recent scenario-based studies have projected billion-dollar-scale future lifecycle or replacement burdens for infrastructure along the QTEC and QTR [13,14]. These estimates differ in spatial scope, price basis, discounting, infrastructure inventory, and adaptation assumptions. They establish the importance of the engineering problem, but they cannot be used to convert a small improvement in ALT estimation into proportional monetary savings. Moreover, ALT is only one component of thaw-settlement hazard: the engineering response to active-layer change also depends on ground-ice content, thaw compressibility, drainage, loading, and asset condition [8,9,11].

Air temperature is the dominant climatic control on permafrost thermal state, but precipitation can modify active-layer dynamics through several opposing pathways. Additional water changes soil thermal conductivity, heat capacity, latent-heat demand, evapotranspiration, and near-surface hydrological connectivity [15–18]. Precipitation can also promote vegetation growth; associated changes in canopy shading, litter and organic-layer properties, snow trapping, and evapotranspiration alter surface energy partitioning and may dampen, enhance, or redistribute heat transfer into the ground [18,19]. Consequently, a precipitation predictor need not have a spatially uniform causal sign. Along the QTEC, precipitation is strongly seasonal: approximately 80% was previously reported to fall during May–September [19], and the site-extracted TPMFD climatology used here assigns 87.1% of annual precipitation to May–September and 63.2% to June–August. These months overlap with active-layer development, vegetation activity, and strong surface–atmosphere exchange.

Previous studies have shown that non-temperature environmental conditions modulate Tibetan Plateau permafrost and that borehole records can retain antecedent climatic information [20,21]. These findings motivate, but do not answer, the present question. The remaining gap is whether seasonally organized precipitation contains reproducible out-of-sample information after stable site differences, long-term change, previous-year ALT, and air temperature have already been considered. Addressing that gap requires more than a full-sample association: candidate windows must be selected without access to test outcomes, uncertainty must respect repeated observations within boreholes, temporal-null tests [22] must retain serial structure, and monitored-site temporal updating must be distinguished from transfer to geographically excluded sites.

The magnitude of any increment is as important as its statistical reproducibility. A small reduction in aggregate prediction error can persist across held-out years without

representing millimeter-scale physical resolution at an individual borehole-year or a material change in an engineering decision. Measurement scale, candidate-selection optimism, dependence among sites and years, and the availability of target-year information must therefore be considered together.

We treated predictive improvement, physical interpretation, geographic transfer, and engineering value as four separate questions. The objectives were to (1) construct annual, warm-season, core-monsoon, and cold-season precipitation windows from monthly TPMFD data; (2) compare post-screening fixed candidates with a pipeline that repeated window selection within each training period; (3) evaluate temporal alignment using corrected, selection-aware circular-shift tests; (4) examine sensitivity to season boundaries, predictor cutoff, rolling start year, learner, and permafrost thermal state; and (5) test strict spatial-block $\times$ time transfer. This hierarchy estimates the magnitude and domain of a predictive increment. It does not identify a unique hydrothermal mechanism, demonstrate a pre-season forecast, or quantify infrastructure maintenance benefits.

2. Materials and Methods

2.1. Study Area, Borehole Observations, and Metadata Screening

The study area covers the Xidatan–Anduo section of the QTEC on the central Tibetan Plateau, extending from approximately 91.53 to 94.08° E and from 32.31 to 35.72° N [6,10]. This section traverses pronounced gradients in latitude, elevation, vegetation, geomorphic setting, and hydrothermal conditions. The 54 boreholes retained for analysis range from 4464 to 5059 m a.s.l. and therefore constitute a heterogeneous monitored transect rather than a single local cluster [20]. For map-based site screening, 20 km buffers [6,23] were constructed around the QTH and QTR alignments between Xidatan and Anduo and dissolved into a single analysis boundary of approximately 24,094 km². All 54 retained boreholes are located within this boundary, as shown in Figure 1.




Figure 1. Study area and locations of the 54 retained boreholes along the Xidatan–Anduo section of the Qinghai–Tibet Engineering Corridor. The analysis boundary was defined using 20 km buffers around the Qinghai–Tibet Highway and Qinghai–Tibet Railway alignments.

Permafrost observations were obtained from the public borehole dataset described by Fu et al. in 2025 [20]. The principal response variable was annual ALT for 2001–2020. The ALT, TTOP, PT10_m, and PT15_m sheets in the released workbook each contained 55 borehole identifiers, whereas the accompanying SITE metadata sheet contained 54. WL3 was the only identifier present in the four observation sheets without a corresponding SITE record; it therefore lacked the coordinates and site attributes required for climate extraction and spatial analysis. In addition, the WL3 columns contained no valid ALT or TTOP values, although PT10_m and PT15_m values were present. WL3 consequently lacked both the ALT response and the spatial metadata required for the ALT–climate analysis and was excluded before model construction. The resulting analytical panel comprised 54 boreholes with complete annual ALT records for all 20 years, yielding 1080 site-year observations.

Observed ALT ranged from 0.82 to 7.80 m, with a median of 2.755 m. These values characterize the monitored panel and should not be interpreted as a spatially uniform thaw regime. TTOP and the annual mean ground-temperature fields at 10 and 15 m depth (source fields PT10_m and PT15_m) were retained for contextual interpretation and cross-variable consistency checks; they were not included as predictors in the primary rolling-origin ALT models. Source missing-value symbols were converted to missing values, and no interpolation or model-based imputation was introduced in the present study. According to the accompanying documentation, ALT was interpreted from temperature–depth profiles with a typical uncertainty of approximately ± 0.05 m [20,21], and source-level quality control addressed outliers, instrument errors, and data gaps. This measurement uncertainty remains part of the response data and is not removed by subsequent statistical modeling.

Available site metadata included longitude, latitude, elevation, ecological type, vegetation cover, volumetric moisture content, source-provided MAGT, geomorphic unit, and soil type. Because MAGT was subsequently used for thermal-state effect-modification and engineering-classification analyses, its consistency with the annual ground-temperature and ALT records was examined before those analyses [15,24]. The released files did not document the measurement depth or reference period represented by the static MAGT field. TG2 was the only site with a positive static MAGT and therefore the only site for which that omission created a material thermal-state provenance question. The source value was retained without correction. TG2 remained in the 54-site primary prediction panel because static MAGT was not part of the primary rolling-origin predictor set, but it was excluded from the principal thermal-state interpretation. The source-record evidence and analysis rule are specified in Section 2.8, and the corresponding 54-site/53-site prediction sensitivity is reported in Supplementary Table S1. Table 1 summarizes the study coverage, sample construction, response uncertainty, validation period, and principal interpretation boundaries.

Table 1. Study data, sample construction, and interpretation boundaries.

Item	Description
Study area	Xidatan–Anduo section of QTEC
Spatial range	32.31–35.72° N; 91.53–94.08° E; 4464–5059 m a.s.l.
Source workbook screening	55 identifiers in each annual observation sheet; WL3 lacked SITE metadata and valid ALT/TTOP and was excluded before climate matching and modeling
Primary ALT panel	54 boreholes $\times$ 20 years = 1080 complete annual ALT records
Thermal-state panel	53 boreholes; TG2 excluded only from the principal thermal-state interpretation under the provenance rule in Section 2.8
TG2 prediction sensitivity	Core 54-site/53-site comparison reported in Supplementary Table S1
ALT period and range	2001–2020; 0.82–7.80 m
ALT source uncertainty	Typical ± 0.05 m from interpreted temperature–depth profiles
Monthly TPMFD	1986–2020; 22,680 complete site-month records
Site-mean annual precipitation	426.6–713.8 mm yr ^{−1} over 1986–2020
Rolling test period	2012–2020; first fold trained on 2002–2011 complete model years
Prediction timing	Retrospective/end-of-season annual estimation, not a pre-season forecast

2.2. Monthly TPMFD and Regional Precipitation Climatology

Monthly near-surface air temperature and precipitation were sampled at the 54 bore-hole coordinates from TPMFD [25,26] for 1986–2020, yielding a complete panel of 22,680 site-month records (54 sites $\times$ 35 years $\times$ 12 months). The archived extraction script selected the nearest TPMFD grid cell to each borehole coordinate; no spatial interpolation between grid cells was applied. Temperature was converted from kelvin to degrees Celsius, and the annual air-temperature variable used in the models was the arithmetic mean of the 12 extracted monthly values. The extracted precipitation field is a monthly mean rate in mm h^{−1}. Monthly total precipitation was calculated as

$$P_m = prcp_m \times 24 \times D_m \quad (1)$$

where P_m is the monthly total (mm), $prcp_m$ is the monthly mean precipitation rate (mm h^{−1}), and D_m is the number of calendar days in month m . Aggregated annual totals reconstructed from the monthly panel agreed with the independently prepared annual TPMFD table within floating-point precision, providing a consistency check on unit conversion and temporal aggregation.

For climatological justification, seasonal fractions were calculated for every complete site-year during 1986–2020 and then summarized across sites. Site-mean annual precipitation ranged from 426.6 to 713.8 mm yr^{−1}. May–September contributed 87.1% of the all-site-year total precipitation; mean site climatological fractions ranged from 72.2% to 91.2%. June–August contributed 63.2%, with site climatological means ranging from 50.4% to 67.2%. These data, together with regional climatological literature [19,27–29], support May–September as the broad thaw/warm season and June–August as a compact core-monsoon window. Because monsoon onset and withdrawal vary, adjacent boundaries were treated as sensitivities rather than independent mechanistic discoveries. The complete climatological summary is reported in Supplementary Table S2a.

2.3. Seasonal and Multi-Year Precipitation Candidates

Four principal seasonal totals were constructed: January–December annual precipitation, May–September warm-season precipitation, June–August core-monsoon precipitation, and the preceding October through current April cold-season precipitation. The cold season was assigned to its ending year, and Table 2 summarizes these four principal definitions. June–September, May–October, preceding November–current March, and preceding October–current May were evaluated separately as boundary sensitivities.

Table 2. Principal seasonal precipitation definitions.

Window	Main Definition	Year Assignment
Annual	January–December	Current year
Warm season	May–September	Current year
Core monsoon	June–August	Current year
Cold season	Previous October–current April	Ending year

For site i , target year t , season s , and a four-year block beginning at lag a , the multi-year seasonal-precipitation predictor was defined as

$$\bar{P}_{i,t,s}^{a:a+3} = \frac{1}{4} \sum_{j=a}^{a+3} P_{i,t-j,s} \quad (2)$$

where $P_{i,t-j,s}$ is the total precipitation for season s at lag j . Four lag blocks (0–3, 4–7, 8–11, and 12–15 years) crossed with the four principal seasonal definitions to create 16 fixed

candidates. Because predictors were standardized within each training fold, the four-year mean rather than the corresponding sum changes only scale, not fitted predictions. The 0–3-year block spans the target year and its three preceding years. The monsoon version reported as the best fixed candidate was identified after comparing all 16 out-of-sample candidate results and is therefore a post-screening descriptive estimate. A separate pipeline selected one of the sixteen candidates using only the training data within each rolling fold.

The annual, warm-season, and monsoon 0–3-year windows include precipitation from target year t . Exact within-year ALT observation dates were unavailable in the public data product. Rolling-origin fitting prevents future years from entering model fitting, but it does not create an operational lead time within year t . Accordingly, the main analysis is interpreted as retrospective or end-of-season annual estimation, not a pre-season forecast. A stricter sensitivity used only June–August precipitation from years $t - 1$ to $t - 4$.

2.4. Rolling-Origin Linear Models

The external test years were 2012–2020. With previous-year ALT required, the first origin used 10 complete model years (2002–2011; 540 site-year records) for training and retained nine consecutive annual test sets, following the principle of rolling-origin time-series evaluation [30,31]. Starting in 2011 or 2013, we examined the sensitivity. At every origin T , only years before T were used for fitting, standardization, temperature-lag selection, precipitation-window selection, and hyperparameter selection. Each test fold contained 54 boreholes, giving 486 held-out site-year predictions across the principal evaluation. Test-year ALT values were used only after predictions had been generated; they did not enter any preprocessing, model selection, fitting, or tuning step.

The baseline model was calculated as

$$ALT_{i,t} = \alpha_i + f(t) + \beta_1 ALT_{i,t-1} + \beta_2 Ta_{i,t-k} + \varepsilon_{i,t} \quad (3)$$

where α_i is a site fixed effect, $f(t)$ is a quadratic time trend, $ALT_{i,t-1}$ is previous-year ALT, and $Ta_{i,t-k}$ is annual air temperature at lag k . The lag $k \in [0, 15]$ was selected by Bayesian information criterion using the training set only. Continuous predictors were standardized using training means and standard deviations and then applied unchanged to the test year.

For every precipitation candidate, an otherwise identical model added the standardized precipitation window. Performance was pooled over all 486 external predictions. Incremental gain was calculated as

$$\Delta RMSE = RMSE_{base} - RMSE_{precip} \quad (4)$$

Positive $\Delta RMSE$ indicates lower out-of-sample error after adding precipitation. A fixed candidate and the training-selected pipeline answer different questions: the former estimates the performance of a named window, whereas the latter estimates the attainable gain when the window must be selected without observing test outcomes.

2.5. Cluster Bootstrap and Corrected Temporal-Null Tests

Uncertainty in $\Delta RMSE$ was summarized using 5000 site-cluster bootstrap repetitions [32]. The paired baseline and precipitation-model squared errors were grouped by borehole; boreholes were sampled with replacement, and all nine held-out errors from a sampled borehole were retained. This preserves both the pairing of competing predictions and within-site dependence across test years. Models were not refitted within each bootstrap repetition, so the interval quantifies site-cluster variation in the realized external prediction errors rather than uncertainty from repeating the complete training process.

A 999-repetition within-site circular-shift test evaluated whether the observed temporal alignment exceeded alignments generated by shifting precipitation in time [22]. All 16 candidate columns were shifted jointly within each site so their mutual structure was retained. For a series of length n , a shift k has circular distance $d(k) = \min(k, n - k)$. The conservative null allowed only $d(k) \geq 4$ years. With the 19-year complete series, admissible shifts were 4–15; shifts 16–18 were excluded because their reverse circular distances are only 3–1 years. The full rolling evaluation, including training-only candidate selection where applicable, was repeated for every null realization.

Two selection-aware statistics were evaluated in every null repetition: (1) the maximum external $\Delta RMSE$ among all 16 fixed candidates and (2) the $\Delta RMSE$ of the pipeline that repeated training-only BIC selection within every rolling fold. Candidate-specific, maximum-statistic, and training-selected null summaries are provided in Supplementary Table S3. The one-sided finite-sample p values used were calculated as

$$p = \frac{1 + \sum(\Delta RMSE_{null} \geq \Delta RMSE_{obs})}{1 + N_{perm}} \quad (5)$$

where $\Delta RMSE_{null}$ is the RMSE reduction obtained from each circularly shifted precipitation series, $\Delta RMSE_{obs}$ is the RMSE reduction obtained from the correctly aligned observed series, and N_{perm} is the total number of circular-shift repetitions (999 in this study). The addition of 1 to both the numerator and denominator provides a finite-sample correction and prevents a zero p value.

Independent site-specific shifts preserve each site's marginal distribution and temporal dependence but weaken synchronous cross-site climate structure. We therefore added a common-shift sensitivity in which all sites received the same admissible circular shift. This retains cross-site coherence but yields only 12 distinct shifts, limiting attainable finite-sample significance.

2.6. Seasonal-Boundary, Prediction-Cutoff, and Rolling-Start Sensitivities

Eight unique seasonal definitions were reconstructed directly from the monthly panel: the four principal definitions and four adjacent-boundary alternatives. Each was evaluated with the same baseline specification, rolling origins, sign convention for $\Delta RMSE$, borehole-cluster bootstrap, and corrected circular-distance null. To avoid treating several nearby warm-season calendars as independent confirmations, a maximum statistic was also calculated across the eight displayed boundaries. Predictor availability was examined separately by replacing the target-year-inclusive 0–3-year mean with a June–August mean over years $t - 1$ to $t - 4$. Finally, the rolling evaluation was repeated with the first test years 2011, 2012, and 2013, corresponding to 9, 10, and 11 initial complete training years. These checks address calendar choice, information cutoff, and evaluation-start dependence without changing the primary model after observing their outcomes.

2.7. Algorithm Benchmarks

Algorithm dependence was examined with ridge regression, random forest, and histogram gradient boosting [33,34]. Every learner used the same outer 2012–2020 rolling origins and the same baseline information: quadratic year terms, previous-year ALT, the training-selected air-temperature lag, and site identity for already monitored sites. Each algorithm was fitted both without and with the fixed monsoon 0–3-year precipitation predictor. Hyperparameters were selected using the two latest training-only inner rolling origins, preserving temporal order inside every outer fold.

The prespecified ridge penalties were 0.1, 1, 10, and 100. Random forests used 200 trees and crossed maximum-feature fractions of 0.5 or 1.0 with minimum leaf sizes of 2 or 5.

Histogram gradient boosting used a learning rate of 0.05 and 250 iterations; four combinations of 7 or 15 terminal nodes, minimum leaf sizes of 10 or 20, and L2 regularization of 0, 1, or 5 were evaluated as coded. The identical outer folds and paired feature sets make the within-learner RMSE difference interpretable as the precipitation increment, whereas absolute RMSE differences among learners also reflect their distinct functional forms and regularization.

2.8. Thermal-State Analysis

Before testing thermal-state effect modification, we examined the source fields for TG2 in greater detail. The site has a static source-provided MAGT of +0.66 °C; its ALT was 4.38 m in 2001 and 7.05 m in 2020, mean PT10_m over 2001–2020 was +0.136 °C, and no valid PT15_m series was available. These observations describe an unusually warm and changing ground-thermal setting, but they do not distinguish among very warm or degrading permafrost, possible talik development, and differences in the depth or reference period represented by the static MAGT field. We therefore did not relabel TG2 or treat its MAGT as a confirmed error. All principal ALT-prediction models retained TG2, while the thermal-state diagnostics described below used the 53-site panel. The core rolling-origin model comparisons were also repeated without TG2 to evaluate whether the prediction result depended on this data-provenance decision (Supplementary Table S1).

Potential effect modification was assessed primarily by adding a centered, standardized $MAGT \times precipitation$ interaction to the same rolling model. Because site fixed effects absorb the time-invariant MAGT main effect, the interaction tests whether the precipitation slope varies with thermal state. External RMSE of the interaction model was compared with that of the precipitation-only model. Site-level gain was calculated from each of the borehole's paired held-out baseline and precipitation-model errors before its association with MAGT were assessed using Spearman correlation. An exploratory full-panel coefficient used site-clustered covariance. The continuous interaction was the primary diagnostic; engineering groups were used only for descriptive stratification because sample sizes were unequal and the thresholds were not estimated from these data.

For engineering-oriented description, sites were grouped following GB 50324–2014 [35]: stable ($MAGT < -2.0$ °C), basically stable ($-2.0 \leq MAGT < -1.0$ °C), unstable ($-1.0 \leq MAGT < -0.5$ °C), and extremely unstable ($MAGT \geq -0.5$ °C). Excluding TG2, these groups contained 6, 15, 14, and 18 sites. A -1 °C QTR-oriented high-/low-temperature split was secondary. Thresholds were treated as engineering classifications, not universal natural discontinuities.

2.9. Strict Spatial-Block $\times$ Time Transfer

The boreholes were ordered by latitude and assigned once to five contiguous blocks containing 10–11 sites; membership remained fixed across all nine test years. For every target year from 2012 to 2020, one block was held out and the model used only earlier years from the other four blocks. Site fixed effects were not used. Two transfer modes were evaluated. The conditional mode retained $ALT(t - 1)$ from each held-out site and therefore represents geographic coefficient transfer when local monitoring history exists; it is not a cold-start test. A deliberately stringent no-target-history mode removed all target-site ALT outcomes and used only year terms, air temperature, latitude, longitude, elevation, volumetric moisture content, and MAGT. Neither mode represents evidence of successful deployment unless precipitation improves its matched baseline.

2.10. Reproducibility and Analytical Workflow

All analyses were conducted from immutable local input snapshots using explicit random seeds, machine-readable intermediate outputs, and traceable scripts. Each main and

supplementary figure is linked to its plotting dataset, and every numerical table is linked to its source result file. Figure 2 summarizes the analytical workflow and validation hierarchy.




Figure 2. Analytical workflow and validation hierarchy. Target-year precipitation is used only for retrospective annual estimation; algorithm, thermal-state, temporal-null, and strict spatial-transfer branches define the interpretation boundary.

3. Results

3.1. Fixed Precipitation Windows and a Training-Selected Pipeline

The baseline external RMSE for 2012–2020 was 0.203491 m. The post-screening fixed June–August 0–3-year candidate reduced RMSE to 0.199312 m, giving $\Delta\text{RMSE} = 0.004179$ m (2.05%). Annual and May–September candidates produced reductions of 0.003637 m (1.79%) and 0.003585 m (1.76%). The October–April candidate worsened RMSE by 0.000528 m. Table 3 reports the matched RMSE values, relative changes, cluster-bootstrap intervals, and candidate-level temporal-null probabilities; Figure 3 displays the same incremental gains on a millimeter RMSE scale without implying millimeter measurement resolution.

Table 3. Core rolling-origin validation results.

Model	RMSE (m)	ΔRMSE (m)	Relative Reduction	Cluster-Bootstrap 95% Interval (m)	Corrected Sitewise-Shift p
Monsoon (June–August), fixed 0–3 yr	0.199312	0.004179	2.05%	0.001745 to 0.006700	0.001
Annual (January–December), fixed 0–3 yr	0.199854	0.003637	1.79%	0.001993 to 0.005369	0.001
Warm season (May–September), fixed 0–3 yr	0.199906	0.003585	1.76%	0.001892 to 0.005447	0.001
Cold season (October–April), fixed 0–3 yr	0.204020	−0.000528	−0.26%	−0.001346 to 0.000343	0.989
Training-selected candidate in each fold	0.201847	0.001645	0.81%	−0.000450 to 0.003948	0.004

Note: The fixed-window estimates are descriptive after candidate screening. The maximum across all 16 fixed candidates had a selection-aware sitewise-shift p value of 0.001. Intervals are percentile site-cluster bootstrap intervals from 5000 resamples of paired external errors; the models were not refitted within each bootstrap repetition. Positive ΔRMSE denotes lower RMSE after adding precipitation.




Figure 3. Incremental retrospective ALT-estimation performance across 486 held-out site-years (2012–2020). Bars show pooled $\Delta RMSE = RMSE_{base} - RMSE_{precip}$; positive values indicate lower error after adding precipitation. Error bars are percentile 95% intervals from 5000 borehole-cluster bootstrap samples of paired external errors. Fixed candidates are descriptive after screening, whereas the training-selected pipeline repeats candidate selection without test outcomes. Millimeter units describe changes in pooled RMSE, not ALT measurement resolution.

The fully training-selected pipeline had $RMSE = 0.201847$ m and $\Delta RMSE = 0.001645$ m (0.81%). Its site-cluster bootstrap interval crossed zero (-0.000450 to 0.003948 m), whereas intervals for the three positive fixed 0–3-year windows did not. The difference between 4.179 mm for the post-screening fixed best candidate and 1.645 mm for the training-selected pipeline quantifies optimism that is relevant when the window must be selected prospectively. These are aggregate error differences; they do not imply that ALT is physically resolved to 1–4 mm, especially because the source reports approximately ± 0.05 m uncertainty for an individual ALT determination.

3.2. Selection-Aware Temporal-Null Evidence

Under corrected sitewise shifts with circular distance ≥ 4 years, the observed maximum across all 16 fixed candidates was 0.004179 m, whereas the 97.5th percentile of the maximum-null distribution was 0.001336 m ($p = 0.001$). The training-selected pipeline had observed $\Delta RMSE = 0.001645$ m, compared with a null 97.5th percentile of 0.001112 m ($p = 0.004$). Thus, the temporal signal remained supported after repeating candidate selection under the stated circular-distance restriction, although the selection-honest training-selected magnitude was smaller. Figure 4a,b visualizes both null distributions, and Supplementary Table S3 provides the candidate-specific and selection-aware numerical summaries.

The common-shift sensitivity gave a more cautious result. Each fixed positive window exceeded all 12 candidate-specific common-shift values, but the smallest attainable corrected finite-sample p was $1/13 = 0.0769$. For the maximum across 16 candidates, one common shift exceeded the observed maximum, giving $p = 0.154$ (Figure 4c). This sensitivity is underpowered because only 12 admissible phases exist, but it shows that the strength of inference depends on whether synchronous cross-site climate structure is preserved.

We therefore describe the temporal evidence as supportive under the sitewise null, not universally definitive.




Figure 4. Corrected temporal-null tests. (a) Distribution of the maximum $\Delta RMSE$ across 16 fixed candidates in 999 site-specific circular-shift repetitions; (b) corresponding null distribution after repeating training-only selection within each rolling fold; (c) the 12 admissible common shifts that preserve synchronous cross-site precipitation structure. Red lines mark observed statistics and dashed lines in (a,b) mark null 97.5th percentiles. The common-shift test has coarse finite-sample resolution and is interpreted as a sensitivity rather than a contradiction of the sitewise test.

3.3. Boundary, Cutoff, and Rolling-Start Sensitivities

Adjacent warm-season definitions were consistent with the main pattern. June–September, June–August, January–December, May–September, and May–October produced $\Delta RMSE$ values of 0.004208, 0.004179, 0.003637, 0.003585, and 0.003253 m. All had candidate-specific corrected sitewise-shift $p = 0.001$, and the maximum across eight unique seasonal boundaries remained supported ($p = 0.001$). October–April, November–March, and October–May produced negative gains. Full results are in Supplementary Table S4 and Figure S1.

However, the cutoff sensitivity that removed target-year precipitation did not support a pre-season claim. Monsoon precipitation averaged over years from $t - 1$ to $t - 4$ worsened RMSE by 0.000506 m (-0.25% ; bootstrap interval from -0.001285 to 0.000291 m; corrected $p = 0.868$; Supplementary Table S2b). This diagnostic retained the otherwise unchanged baseline, including the air-temperature lag selected from each training set; that lag equaled 0 in the 2018–2020 folds. It is therefore a no-target-year-precipitation sensitivity, not a fully specified pre-season pipeline. Consequently, the target-year windows should not be presented as operational forecasts. The fixed monsoon result was positive when the first test year was 2011, 2012, or 2013, with $\Delta RMSE = 0.003856$, 0.004179 , and 0.004657 m, respectively, showing that the main magnitude was not specific to the 2012 cutoff (Supplementary Table S2c).

3.4. Algorithm Dependence

The paired learner comparison was heterogeneous (Table 4; Figure 5). Adding the fixed monsoon predictor reduced RMSE by 0.004179 m for OLS, 0.003994 m for ridge regression, and 0.004879 m for random forest, but increased RMSE by 0.004524 m for histogram gradient boosting. The precipitation-augmented OLS and ridge models had the lowest absolute RMSE values (0.199312 and 0.199644 m, respectively). Thus, the small increment was not unique to unregularized OLS, but its sign was not invariant to learner choice. These are paired benchmark estimates under matched information and folds, not a formal superiority test proving that one algorithm is generally optimal.

Table 4. Paired model-framework comparison.

Algorithm	Baseline RMSE (m)	With Precipitation RMSE (m)	Δ RMSE (m)
Histogram gradient boosting	0.224500	0.229024	−0.004524
OLS	0.203491	0.199312	0.004179
Random forest	0.215696	0.210817	0.004879
Ridge	0.203638	0.199644	0.003994

Note: Every pair uses the same 2012–2020 outer rolling origins and matched information set. Hyperparameters were selected using training-only inner origins. Positive Δ RMSE denotes improvement within the same learner; the table is not a formal significance ranking of algorithms.




Figure 5. Paired algorithm benchmarks using identical 2012–2020 outer rolling origins, matched feature sets, and training-only inner tuning. Each learner is shown without and with the fixed June–August 0–3-year precipitation predictor. Positive Δ RMSE means improvement relative to the same learner without precipitation; the comparison is descriptive rather than a formal algorithm-ranking test.

3.5. Thermal-State Sensitivity

The 53-site thermal-state analysis did not support a stronger precipitation increment in warmer permafrost. The precipitation-only model had an external RMSE of 0.200040 m, whereas adding the continuous $MAGT \times precipitation$ interaction increased RMSE to 0.200865 m (Δ RMSE relative to the precipitation-only model = −0.000825 m; site-cluster bootstrap interval −0.002018 to 0.000335 m). The exploratory site-clustered interaction coefficient was not statistically distinguishable from zero ($p = 0.356$). Consistently, the site-level association between MAGT and precipitation gain was weak and negative (Spearman $\rho = -0.153$, $p = 0.275$; bootstrap interval −0.394 to 0.110), and fold-specific interaction coefficients were negative in six of nine folds. These results provide no evidence that the precipitation increment increased systematically as permafrost became warmer; they do not prove the absence of all thermal-state modification.

Engineering-class estimates were non-monotonic and imprecise, with the six-site stable group showing a larger descriptive gain than the warmer groups. These results do not justify the hypothesis that precipitation information becomes systematically more valuable in high-temperature permafrost. Excluding TG2 slightly increased the fixed monsoon Δ RMSE from 0.004179 to 0.004486 m, indicating that the aggregate prediction result did

not materially depend on that site's uncertain thermal-state metadata. Detailed class and TG2 sensitivities are reported in Supplementary Tables S1, S5, and S6 and Figure S2.

3.6. Strict Geographic Transfer

Precipitation did not improve strict spatial-block $\times$ time transfer (Table 5; Figure 6). In the conditional mode, which retained held-out-site $ALT(t-1)$ but used no held-out site effect, pooled baseline RMSE was 0.202596 m and precipitation-augmented RMSE was 0.204812 m ($\Delta RMSE = -0.002215$ m). All five held blocks had negative gain; block-specific sample sizes and errors are provided in Supplementary Table S7.

Table 5. Pooled strict spatial-block $\times$ time transfer.

Transfer Mode	Sites	N	Baseline RMSE (m)	With Precipitation RMSE (m)	$\Delta RMSE$ (m)
Conditional transfer with held-site $ALT(t-1)$	54	486	0.202596	0.204812	−0.002215
No target-site ALT history; static metadata only	54	486	1.039164	1.138623	−0.099459

Note: Five fixed latitudinal blocks were held out across nine target years. The conditional mode retains held-site $ALT(t-1)$ but no site fixed effect; the no-target-history mode removes all target-site ALT outcomes. Negative $\Delta RMSE$ denotes deterioration after adding precipitation.




Figure 6. Strict spatial-block $\times$ time holdout across five fixed latitudinal blocks and nine target years. Every target year was predicted from earlier years in the other four blocks: (a) retains held-site $ALT(t-1)$ but no site fixed effect; (b) removes all target-site ALT history. Positive $\Delta RMSE$ would indicate improvement. Pooled and block-specific results show no transferable precipitation gain.

Without any target-site ALT history, the static-covariate baseline had $RMSE = 1.039164$ m and the precipitation model had $RMSE = 1.138623$ m ($\Delta RMSE = -0.099459$ m). This deliberately stringent diagnostic performed poorly in absolute and relative terms. The fitted relationship is therefore restricted to temporal updating where a site's own monitoring history is available; it is not supported for cold-start prediction at never-monitored corridor locations.

4. Discussion

4.1. Contribution and Magnitude

Recent Tibetan Plateau studies have shown that non-temperature environmental conditions can modulate warming-induced permafrost degradation and that borehole

records retain decadal-scale climatic memory [20,21]. Those findings establish that ground conditions reflect more than contemporaneous air temperature, but they do not determine whether monthly precipitation can be reorganized into seasonal windows that add out-of-sample information after local baseline behavior and antecedent ALT have been considered. The present study addresses that narrower question. It is therefore a seasonal incremental-prediction analysis rather than a broad attribution analysis of permafrost degradation.

The main contribution is not a large improvement in ALT estimation. It is an explicit estimate of the small information increment associated with precipitation timing after stable site differences; temporal trend, previous-year ALT, and temperature have already been considered. This is a deliberately conditional quantity: it asks whether precipitation reduces the residual predictive error of a strong monitored-site baseline, not whether precipitation is a dominant control of ALT.

The post-screening and training-selected results describe different estimands. The strongest fixed candidate reduced pooled RMSE by 4.179 mm after the 16 candidates had been compared, whereas the pipeline that had to choose its window from training data alone reduced it by 1.645 mm. The difference provides an empirical indication of selection optimism. Selection-aware sitewise temporal-null tests remained supportive, but the common-shift test preserving synchronous cross-site climate structure was inferentially inconclusive. Site-cluster bootstrap intervals additionally describe variation in the realized paired errors across boreholes; because models were not refitted in each bootstrap repetition, they should not be interpreted as uncertainty for the entire model-development process. This evidence hierarchy is more informative than describing the result simply as robust.

The magnitude also needs to be interpreted against measurement scale. The source-reported ± 0.05 m uncertainty is more than ten times the fixed-window RMSE reduction and approximately thirty times the training-selected reduction. Aggregate RMSE differences can be reproducible across 486 predictions even when they are smaller than individual-observation uncertainty because RMSE pools information over sites and years. Nevertheless, the gain cannot be interpreted as millimeter-scale physical measurement accuracy, a directly observable change at a single site-year, or evidence that the model corrects source-level measurement error. Response uncertainty may also attenuate weak climatic relationships and limit the practical distinguishability of competing predictions.

4.2. Physical Interpretation of Warm-Season Timing

The warm-season timing is process-consistent but does not constitute physical proof. During thaw, rainfall can increase liquid-water availability and soil thermal conductivity, while additional water also increases heat capacity and latent-heat demand. Drainage, soil texture, ground ice, and surface water determine whether water promotes downward heat transfer or buffers temperature change [15–18]. Precipitation-driven vegetation growth can further modify shading, litter, roughness, evapotranspiration, and snow trapping [18,19]. These pathways can produce different net effects among sites.

These competing pathways also mean that the sign of a fitted precipitation coefficient is conditional on the selected covariates, temporal window, and training sample. The present design estimates predictive increment rather than a causal dose–response relationship; predictive gain does not establish that more precipitation increases ALT or engineering risk. The lack of evidence for systematic MAGT moderation further argues against a simple, warmer permafrost–stronger rainfall effect explanation.

The annual candidate likely performs well because it contains the dominant warm-season component. The local TPMFD climatology assigns 87.1% of annual precipitation to May–September and 63.2% to June–August, while neighboring warm-season boundaries

yield similar gains. The evidence therefore points to information concentrated in the thaw-season hydroclimatic state, not equal contributions from all months.

4.3. Possible Explanations for the Lack of Cold-Season Predictive Gain

The absence of a cold-season gain was consistent across alternative calendar definitions rather than being an artifact of one boundary. The preceding October–current April, preceding November–current March, and preceding October–current May totals all produced negative $\Delta RMSE$ values; for the October–May definition, the borehole-cluster interval was also entirely below zero. Thus, none of the three cold-season totals added predictive information to the matched temperature-and-persistence baseline (Table 3; Supplementary Table S4; Figure S1). This convergence across adjacent calendars is more informative than the result from any single boundary.

This predictive result does not imply that winter snow is thermally unimportant. The surface insulation relevant to frozen-ground temperature depends on snow depth, density, thermal conductivity, stratigraphy, duration, onset, melt timing, sublimation, and liquid-water refreezing, not simply on accumulated precipitation [36,37]. Rain–snow partitioning can also vary within a nominal cold season. In the windy alpine corridor, preferential deposition, scouring, and redistribution may create large differences in snow cover over short distances even when grid-cell precipitation totals are similar [36]. Consequently, TPMFD cold-season precipitation is an incomplete proxy for the snowpack properties that control surface thermal resistance.

The conditional model structure provides a second explanation for the small marginal contribution. Annual air temperature represents a major component of atmospheric forcing, while $ALT(t - 1)$ may proxy integrated information from the previous surface energy balance, soil moisture, snow regime, and ground thermal state. If cold-season precipitation affects ALT mainly through processes already expressed in these variables, little independent information will remain after conditioning. In addition, the spatial scale of gridded forcing may be mismatched to borehole-scale snow redistribution, and a seasonal sum cannot resolve whether snowfall occurred early enough to insulate the ground, late enough to delay thaw, or during transient melt–refreeze events.

The appropriate interpretation is therefore conditional and scale specific: aggregated cold-season precipitation did not improve annual ALT estimation in the present data and model, but this is not evidence against snow insulation as a physical control or against its relevance to engineering design. Direct observations of snow depth, snow–water equivalent, density, accumulation and disappearance dates, wind redistribution, surface temperature, and freeze–thaw energy fluxes would be needed to test the winter mechanism. Until such observations are available, the negative predictive result should be used to delimit the current covariate set, not to dismiss winter surface processes.

4.4. Algorithm and Model-Selection Dependence

Ridge regression and random forest reproduced a small paired precipitation gain, whereas histogram gradient boosting did not. Complex learners do not automatically improve a short panel with strong site persistence and only nine external annual origins. The linear model encodes the dominant structure explicitly through site effects, temporal trend, autoregression, and temperature. Ridge regression retains that structure while shrinking coefficients, and random forest can represent limited nonlinearities and interactions [33]. Histogram gradient boosting partitions the feature space more aggressively; with correlated predictors and a modest number of temporally independent test origins, that flexibility may increase estimation variance or allocate capacity to patterns that do not persist out of sample.

The negative gradient-boosting result is therefore scientifically useful, but it is not proof that nonlinear hydrothermal responses are absent. It shows that the tested precipitation increment was not invariant to learner and tuning grid under the available sample. Conversely, positive paired differences for OLS, ridge, and random forest do not constitute three independent significance tests because the learners used the same observations, folds, and precipitation predictor. OLS remains a defensible primary framework because it had the lowest augmented-model RMSE and makes the conditioning structure transparent, not because the benchmark establishes universal linear superiority.

The distinction between fixed and selected candidates is equally important. The 4.179 mm value describes the best named candidate after broad comparison; the 1.645 mm value better represents a model that must choose a window without test outcomes. Both are reported, rather than retrospectively labeling the best candidate as prespecified.

4.5. Temporal Updating Versus Geographic Transfer

Temporal and geographic validation answer different questions and can legitimately produce different conclusions. Temporal evaluation asks whether precipitation history improves retrospective estimation after a monitored site's own persistence, baseline level, and previous ALT have been learned. Strict geographic validation removes the held block's site effects and asks whether relationships estimated from other corridor blocks transfer into later years [21]. The conditional spatial test still supplies $ALT(t - 1)$ for the held site, yet precipitation worsened performance in all five blocks. This indicates that the small within-site temporal increment did not translate into a stable cross-block coefficient adjustment.

The no-target-history mode is more stringent. Its large absolute error reflects the loss of site-specific ALT level and persistence as well as limited transferability of the available static covariates; it should not be attributed to precipitation alone. The paired comparison nevertheless remains informative because adding precipitation also worsened that matched baseline. Together, the two modes separate, local history is known but coefficients are imported from no local ALT outcome has ever been observed, avoiding the common mistake of treating a held-site lag as evidence for ungauged deployment.

Weak geographic transfer narrows the domain of use. Local ground ice, soil, vegetation, topography, drainage, disturbance, and engineering structure mediate precipitation-ground interactions [7–9,20,26,29,38]. These properties are not adequately represented by a single gridded precipitation term. Prediction at never-monitored locations requires transferable subsurface and surface covariates and independent validation at sites excluded from every stage of model development.

4.6. Engineering Relevance and Economic Boundary

The engineering relevance is limited but identifiable. At a site with an established ALT record, seasonal precipitation history can be added to a multi-source annual assessment to quantify whether a temperature-and-persistence baseline is less adequate in a particular year [23]. Whether that information could help prioritize data review is a value-of-information hypothesis that requires prospective decision records; it has not been demonstrated here. The immediate implication is instead to motivate closer integration with direct thermal and deformation observations. However, the no-target-year-precipitation sensitivity showed no gain, so the present model is not an advance-warning trigger. Nor should precipitation be used as a design parameter, stand-alone hazard index, or substitute for deformation, ground-temperature, and ground-ice observations.

This study contains no infrastructure deformation, damage, inspection, intervention, maintenance, or cost outcomes. It therefore cannot show that a 0.81–2.05% RMSE reduction yields measurable maintenance savings. The historical and scenario-based economic

literature [11–13] establishes that permafrost-related infrastructure impacts can be large, but it does not monetize this model’s incremental information. A value-of-information analysis would require asset-specific deformation, ground ice, damage thresholds, inspection decisions, treatment costs, avoided losses, and evidence that the estimate changes a decision.

The defensible practical conclusion is therefore modest: seasonal precipitation is a contextual covariate for retrospective assessment at already monitored sites. Its potential value lies in enriching a broader evidence system, not in independently triggering engineering intervention.

4.7. Limitations and Future Work

Several limitations remain. First, exact within-year ALT observation dates were unavailable, and the main windows contain target-year climate. The model is retrospective; the fully prior-year precipitation sensitivity did not improve prediction. Second, the best fixed gain is much smaller than individual ALT uncertainty and was selected after candidate comparison. Although maximum-statistic and training-selected null tests were supportive, a common-shift null preserving cross-site coherence was inconclusive because only 12 phases were available.

Third, TPMFD is gridded. Microtopography, vegetation patches, drainage, engineering disturbance, and wind-driven snow redistribution can make borehole conditions differ from grid-cell climate [22,25,34,36]. Fourth, no soil moisture time series, snow measurements, surface-energy fluxes, or ground-ice profiles were available to isolate mechanisms. Coefficient signs from these conditional predictive models must not be generalized causally.

Fifth, the released static MAGT field lacks a documented measurement depth and reference period; this limitation is particularly consequential for the positive value at TG2 and prevents an unambiguous classification of that site. Thermal-state inference was therefore based on 53 sites, while the 54-site/53-site comparison showed that the core prediction result was not dependent on TG2. The stable engineering class contained only six sites, limiting subgroup precision. Sixth, the 2012–2020 evaluation provides nine annual test origins rather than a long independent record. Finally, strict spatial transfer failed, and no infrastructure-response or cost data were available.

Future work should collect observation timing, snow depth and snow-water equivalent, soil moisture, vegetation dynamics, ground-ice content, deformation, and asset-condition data. Prospective models should define an explicit predictor cutoff and lead time. Spatial models should be validated at fully excluded sites and coupled to physical and engineering covariates. Economic value should be assessed only after model output can be linked to real decisions and observed outcomes.

5. Conclusions

- (1) After site effects, quadratic trend, previous-year ALT, and a training-selected air-temperature lag were considered, recent precipitation contained a small increment of information for retrospective ALT estimation at monitored sites. The post-screening June–August 0–3-year candidate reduced RMSE by 0.004179 m (2.05%), whereas a training-selected pipeline reduced it by 0.001645 m (0.81%).
- (2) Corrected, selection-aware sitewise circular-shift tests supported both statistics, but a common-shift sensitivity preserving cross-site climate coherence was inconclusive. Adjacent warm-season definitions were consistent, whereas cold-season totals were unsupported. A predictor using only years $t - 1$ to $t - 4$ did not improve performance; the study therefore does not demonstrate a pre-season forecast.

- (3) Ridge regression and random forest showed small paired gains, but histogram gradient boosting did not. A continuous MAGT interaction provided no evidence of stronger precipitation information in warmer permafrost. Strict spatial-block $\times$ time holdout showed no transferable gain, even when held-out-site previous-year ALT was available.
- (4) The RMSE increment is much smaller than individual ALT uncertainty and cannot be interpreted as millimeter-scale physical resolution. The study contains no infrastructure damage, maintenance, or cost outcomes. The supported use is limited to an auxiliary covariate in retrospective multi-source assessment at sites with established monitoring histories; it does not support ungauged-site hazard mapping or quantified economic benefit.

Supplementary Materials: The following supporting information can be downloaded at <https://www.mdpi.com/article/10.3390/buildings16153023/s1>: Figure S1: Seasonal-boundary sensitivity; Figure S2: Thermal-state sensitivity; Table S1: TG2 source-record and inclusion/exclusion sensitivity; Table S2: Climatology, rolling-start, and no-target-year-precipitation sensitivities; Table S3: Corrected temporal-null summary; Table S4: Corrected seasonal-boundary results; Table S5: Engineering thermal-state groups; Table S6: Continuous thermal interaction; Table S7: Strict spatial-block results.

Author Contributions: Conceptualization, Q.D. and F.W.; methodology, Q.D., F.W. and G.L.; software, Q.D., F.W. and S.Q.; validation, Q.D., F.W., G.L. and D.C.; formal analysis, Q.D., D.C. and S.Q.; investigation, Q.D., F.W., D.C. and S.Q.; resources, Q.D.; data curation, Q.D. and F.W.; writing—original draft preparation, Q.D. and F.W.; writing—review and editing, Q.D., F.W. and G.L.; visualization, Q.D. and S.Q.; funding acquisition, Q.D. and G.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Program of the Science and Technology Program of Xizang Autonomous Region, grant number XZ202401ZY0040, the Program of the Gansu Province Science and Technology Foundation for Youths, grant number 25JRR515, the Research Project of the Qinghai Provincial Key Laboratory of Tibet Plateau Highway Construction and Maintenance Technology, grant number 2024-JY-D-03, and the Research Project of the Heilongjiang Provincial Hydraulic Research Institute, grant number DT2024B01.

Data Availability Statement: The borehole permafrost observation dataset is available from Figshare at <https://doi.org/10.6084/m9.figshare.29206613.v1> (accessed on 20 June 2026). The climate forcing data are available from the Third Pole Meteorological Forcing Dataset at <https://cstr.cn/18406.11.Atmos.tpd.c.302088> (accessed on 20 June 2026). Other related and processed datasets are provided in the Supplementary Materials.

Acknowledgments: The authors acknowledge the public availability of the borehole observation data and the Third Pole Meteorological Forcing Dataset. The authors are also grateful to the University of Alberta for providing access to ArcGIS Pro 3.6.0, which supported data processing and figure generation in this study.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Yang, S.; Wen, X.; Wu, T.; Wu, X.; Wang, X.; Jin, X.; Li, X.; Yang, X.; Yang, L.; Wang, H. Carbon-Cycling Microorganisms in Permafrost and Their Responses to a Warming Climate: A Review. *Permafr. Periglac. Process.* **2024**, *35*, 218–231. [CrossRef]
2. Wang, Y.; Gao, H.; Jin, H.; Xi, Q.; Wang, P.; Chen, D. Root Zone Storage Capacity Reveals Ecohydrological Turning Points in Tibetan Plateau Permafrost Regions. *Catena* **2026**, *263*, 109716. [CrossRef]
3. Ji, F.; Fan, L.; Kuang, X.; Li, X.; Cao, B.; Cheng, G.; Yao, Y.; Zheng, C. How Does Soil Water Content Influence Permafrost Evolution on the Qinghai-Tibet Plateau under Climate Warming? *Environ. Res. Lett.* **2022**, *17*, 064012. [CrossRef]
4. Zhou, L.; Yang, Y.; Zhang, D.; Yao, H. Recent Advances in Hydrology Studies under Changing Permafrost on the Qinghai-Xizang Plateau. *Res. Cold Arid Reg.* **2024**, *16*, 159–169. [CrossRef]

5. Du, Q.; Li, G.; Chen, D.; Zhou, Y.; Qi, S.; Wang, F.; Mao, Y.; Zhang, J.; Cao, Y.; Gao, K.; et al. Bibliometric Analysis of the Permafrost Research: Developments, Impacts, and Trends. *Remote Sens.* **2023**, *15*, 234. [\[CrossRef\]](#)
6. Du, Q.; Chen, D.; Li, G.; Cao, Y.; Zhou, Y.; Chai, M.; Wang, F.; Qi, S.; Wu, G.; Gao, K.; et al. Preliminary Study on InSAR-Based Uplift or Subsidence Monitoring and Stability Evaluation of Ground Surface in the Permafrost Zone of the Qinghai–Tibet Engineering Corridor, China. *Remote Sens.* **2023**, *15*, 3728. [\[CrossRef\]](#)
7. Sun, Z.; Zhao, L.; Hu, G.; Zhou, H.; Liu, S.; Qiao, Y.; Du, E.; Zou, D.; Xie, C. Numerical Simulation of Thaw Settlement and Permafrost Changes at Three Sites Along the Qinghai–Tibet Engineering Corridor in a Warming Climate. *Geophys. Res. Lett.* **2022**, *49*, e2021GL097334. [\[CrossRef\]](#)
8. Ni, J.; Wu, T.; Zhu, X.; Wu, X.; Pang, Q.; Zou, D.; Chen, J.; Li, R.; Hu, G.; Du, Y.; et al. Risk Assessment of Potential Thaw Settlement Hazard in the Permafrost Regions of Qinghai–Tibet Plateau. *Sci. Total Environ.* **2021**, *776*, 145855. [\[CrossRef\]](#) [\[PubMed\]](#)
9. Liu, Z.; Zhu, Y.; Chen, J.; Cui, F.; Zhu, W.; Liu, J.; Yu, H. Risk Zoning of Permafrost Thaw Settlement in the Qinghai–Tibet Engineering Corridor. *Remote Sens.* **2023**, *15*, 3913. [\[CrossRef\]](#)
10. Du, Q.; Li, G.; Che, T.; Mu, Y.; Ma, W. A Dataset of InSAR-Derived Ground Deformation in Permafrost Zone of the Qinghai–Tibet Engineering Corridor, China (2017–2022). *China Sci. Data* **2025**, *10*, 1–20. [\[CrossRef\]](#)
11. Ruan, G.; Zhang, J.; Chai, M. Risk division of thaw settlement hazard along Qinghai–Tibet Engineering Corridor under climate change. *J. Glaciol. Geocryol.* **2014**, *36*, 811–817.
12. Ma, W.; Mu, Y.; Xie, S.; Mao, Y.; Chen, D. Thermal-Mechanical Influences and Environmental Effects of Expressway Construction on the Qinghai–Tibet Permafrost Engineering Corridor. *Adv. Earth Sci.* **2017**, *32*, 459. [\[CrossRef\]](#)
13. Ni, W.-H.; Yin, G.-A.; Niu, F.-J.; Lin, Z.-J.; Luo, J.; Yan, H.-Y.; Gao, Z.-Y.; Ju, X.; Liu, Q. High-Resolution Risk Assessment Reveals Billion-Dollar Additional Threat to Infrastructure from Permafrost Thaw in the Qinghai–Tibet Engineering Corridor. *Adv. Clim. Change Res.* **2026**, *17*, 615–628. [\[CrossRef\]](#)
14. Zhang, Z.; Chen, F.; Lin, H.; Wang, C.; Liu, X.; Wang, M.; Luo, J. Satellite Observations Characterize the Impacts of Climate Change and Human Activities on Permafrost along Qinghai–Tibet Railway. *Innov. Geosci.* **2025**, *3*, 100127. [\[CrossRef\]](#)
15. Zhang, M.; Wen, Z.; Li, D.; Chou, Y.; Zhou, Z.; Zhou, F.; Lei, B. Impact Process and Mechanism of Summertime Rainfall on Thermal-Moisture Regime of Active Layer in Permafrost Regions of Central Qinghai–Tibet Plateau. *Sci. Total Environ.* **2021**, *796*, 148970. [\[CrossRef\]](#) [\[PubMed\]](#)
16. Muir, G.; Brown, G.S.; Balasubramaniam, K.; Hu, B. Active Layer Thermal Regime Varies across Landforms in a Subarctic Wetland. *Facets* **2025**, *10*, 1–14. [\[CrossRef\]](#)
17. Xu, X.; Wu, Q. Active Layer Thickness Variation on the Qinghai–Tibetan Plateau: Historical and Projected Trends. *J. Geophys. Res.-Atmos.* **2021**, *126*, e2021JD034841. [\[CrossRef\]](#)
18. Xiao, Y.; Hu, G.; Zhao, L.; Du, E.; Li, R.; Wu, T.; Wu, X.; Liu, G.; Zou, D.; Xing, Z.; et al. Vegetation and Permafrost Interactions Shape Soil Moisture Stratification in Marginal Permafrost Zones. *Geoderma* **2025**, *464*, 117596. [\[CrossRef\]](#)
19. Luo, L.; Ma, W.; Zhuang, Y.; Zhang, Y.; Yi, S.; Xu, J.; Long, Y.; Ma, D.; Zhang, Z. The Impacts of Climate Change and Human Activities on Alpine Vegetation and Permafrost in the Qinghai–Tibet Engineering Corridor. *Ecol. Indic.* **2018**, *93*, 24–35. [\[CrossRef\]](#)
20. Fu, Z.; Wu, Q.; Chen, A.; Wang, L.; Jiang, G.; Gao, S.; Yun, H.; Chen, J. Non-Temperature Environmental Drivers Modulate Warming-Induced 21st-Century Permafrost Degradation on the Tibetan Plateau. *Nat. Commun.* **2025**, *16*, 7556. [\[CrossRef\]](#) [\[PubMed\]](#)
21. Fu, Z.; Wang, L.; Jiang, G.; Men, X.; Du, W.; Yang, Y.; Gao, S.; Zhang, Z.; Wu, Q. Decadal-Scale Thermal Memory of Permafrost and Climatic and Topographic Modulation on the Tibetan Plateau. *npj Clim. Atmos. Sci.* **2026**, *9*, 100. [\[CrossRef\]](#)
22. Romano, J.P.; Tirlea, M.A. Permutation Testing for Dependence in Time Series. *J. Time Ser. Anal.* **2022**, *43*, 781–807. [\[CrossRef\]](#)
23. Du, Q.; Xu, A.; Wang, F.; Luo, H.; Qi, S. Analysis of Environmental Driving Mechanisms for Vertical Surface Deformation in the Permafrost Section of the Qinghai–Tibet Engineering Corridor, China. *J. Meas. Eng.* **2026**, *14*, 236–249. [\[CrossRef\]](#)
24. Brown, N.; Gruber, S. Beyond MAGT: Learning More from Permafrost Thermal Monitoring Data with Additional Metrics. *Cryosphere* **2026**, *20*, 1771–1796. [\[CrossRef\]](#)
25. Jiang, Y.; Yang, K.; Qi, Y.; Zhou, X.; He, J.; Lu, H.; Li, X.; Chen, Y.; Li, X.; Zhou, B.; et al. TPhIPr: A Long-Term (1979–2020) High-Accuracy Precipitation Dataset (1/30°, Daily) for the Third Pole Region Based on High-Resolution Atmospheric Modeling and Dense Observations. *Earth Syst. Sci. Data* **2023**, *15*, 621–638. [\[CrossRef\]](#)
26. Jiang, Y.; Tang, W.; Yang, K.; He, J.; Shao, C.; Zhou, X.; Lu, H.; Chen, Y.; Li, X.; Shi, J. Development of a High-Resolution near-Surface Meteorological Forcing Dataset for the Third Pole Region. *Sci. China-Earth Sci.* **2025**, *68*, 1274–1290. [\[CrossRef\]](#)
27. Liu, H.; Liu, X.; Liu, C.; Yun, Y. High-Resolution Regional Climate Modeling of Warm-Season Precipitation over the Tibetan Plateau: Impact of Grid Spacing and Convective Parameterization. *Atmos. Res.* **2023**, *281*, 106498. [\[CrossRef\]](#)
28. Zhang, S.; Meng, L.; Zhao, Y.; Yang, X.; Huang, A. The Influence of the Tibetan Plateau Monsoon on Summer Precipitation in Central Asia. *Front. Earth Sci.* **2022**, *10*, 771104. [\[CrossRef\]](#)
29. Zhang, Y.S.; Ohata, T.; Kadota, T. Land-Surface Hydrological Processes in the Permafrost Region of the Eastern Tibetan Plateau. *J. Hydrol.* **2003**, *283*, 41–56. [\[CrossRef\]](#)

30. Liu, M.; Wang, L.; Sun, L.; Zheng, F. An Improved Feature-Based Forecast Combination Method Using Rolling Origin Evaluation. *Oper. Res. Lett.* **2026**, *68*, 107468. [[CrossRef](#)]
31. Bergmeir, C.; Hyndman, R.J.; Koo, B. A Note on the Validity of Cross-Validation for Evaluating Autoregressive Time Series Prediction. *Comput. Stat. Data Anal.* **2018**, *120*, 70–83. [[CrossRef](#)]
32. Ng, E.S.-W.; Grieve, R.; Carpenter, J.R. Two-Stage Nonparametric Bootstrap Sampling with Shrinkage Correction for Clustered Data. *Stata J.* **2013**, *13*, 141–164. [[CrossRef](#)]
33. Breiman, L. Random Forests. *Mach. Learn.* **2001**, *45*, 5–32. [[CrossRef](#)]
34. Friedman, J.H. Greedy Function Approximation: A Gradient Boosting Machine. *Ann. Stat.* **2001**, *29*, 1189–1232. [[CrossRef](#)]
35. GB 50324–2014; Code for Investigation of Geotechnical Engineering on Frozen Ground. China Planning Press: Beijing, China, 2014.
36. Zhang, T. Influence of the Seasonal Snow Cover on the Ground Thermal Regime: An Overview. *Rev. Geophys.* **2005**, *43*, RG4002. [[CrossRef](#)]
37. Wang, C.; Shirley, I.A.; Wielandt, S.; Fiolleau, S.; Lamb, J.R.; Uhlemann, S.; Ulrich, C.; Bennett, K.E.; Dafflon, B. Advancing the Understanding of Snow Accumulation, Melting, and Associated Thermal Insulation Using Spatially Dense Snow Depth and Temperature Time Series. *Geophys. Res. Lett.* **2025**, *52*, e2024GL114189. [[CrossRef](#)]
38. Fan, X.-W.; Gao, Z.-Y.; Luo, J.; Niu, F.-J.; Li, W.-J.; Wu, X.-Y.; Lin, Z.-J. Spatial Distribution and Zonation Characteristics of Permafrost Ground Ice on the Qinghai–Xizang Engineering Corridor, China. *Adv. Clim. Change Res.* **2026**, *17*, 551–562. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.